LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-590572

Twenty- Five-Year High Performance Computing Facility Master Plan and Recommendation

M. McCoy

October 11, 2012

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Lawrence Livermore National Laboratory

ASC Program Executive Brief:

Twenty-Five-Year High Performance Computing Facility Master Plan and Recommendation for Livermore Computing

FY13–FY14 General Plant Project

M. McCoy

October 2012

LLNL-TR-590572

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

Disclaimer

ASC Program Executive Brief

25 Year High Performance Computing Facility Master Plan and Recommendation for LC FY13-FY14 GPP Project

Executive Summary

The current Terascale Simulation Facility (TSF) Complex facilities at Lawrence Livermore National Laboratory (LLNL) include buildings B-115, B-117, B-439, B-451, and B-453 (the TSF). B-115 and B117 are 60 years old, have permanent limitations with respect to the ability to upgrade their electrical and mechanical infrastructure, and have significant structural deficiencies that must be addressed and are prohibitively costly given the age and limitations of the buildings. They are slated for Deactivation and Decommissioning (D&D).

B-439 and B-451 house most of the current unclassified systems[1] but are inadequate without major infrastructure upgrades to support future (near term future, within 1-3 years) systems including advanced development, commodity, or file systems. However, they are adequate for less dense support servers and storage, and therefore can absorb that type of equipment from buildings B-115 and B117 after some modest upgrades.

Even B-453, the primary LC facility commissioned in 2004, is facing potential problems in terms of supporting future advanced development and commodity systems. Significant electrical, mechanical, and structural upgrades may be necessary depending on the technological advances that are made (or not made) in the future. B-453 also is and will continue to be dedicated primarily to hosting classified systems. Supporting unclassified systems is problematic and reduces flexibility in siting future classified systems but is done when there is no other choice, such as with Vulcan, the 5PF unclassified BGQ system.

By far the most cost effective solution to acquire sufficient unclassified computer room space suitable for supporting the Advanced Simulation and Computing (ASC) Program's mid-term advanced development and commodity systems is to build a modern, scalable and efficient structure employing LLNL's design for a *Sustainable High Performance Computing (HPC) Modular Facility*. The construction cost for a 6,000SF, 6MW module, including mechanical cooling is approximately $10M. The only alternative is to upgrade B-451, at a cost of $20.5M and the program would still be left with an aging facility that will be at its extreme limits in terms of future upgradeability and capacity. The *Sustainable HPC Modular Facility* provides the

[1] Except for Vulcan, the 24-rack unclassified BG/Q. Neither B-439 nor B-451 had the necessary infrastructure to be able to site Vulcan.

program with a flexible infrastructure that can be grown with the addition of modules, but only as needed and only with technologies and requirements that are relevant *at that time*.

Design for the Susta*inable HPC Modular Facility* is complete; however there may be some minor modifications effected to optimally meet changing DOE energy and efficiency requirements. We recommend that the project be funded through GPP, initiating in the second half of FY13 and completing in the first half of FY14.

LLNL requests that the HQ NNSA ASC program concur with execution of the first module of this effort. The LLNL ASC program suggests that $5M FY13 and $5M FY14 of LC FOUS infostructure operating funds be used for a GPP project to complete the module in FY14. LLNL ***does not*** *request additional funding from HQ for this work. The 5-year LLNL infostructure plan has allocated the necessary funding, as this strategy is far less expensive than retrofitting B451, the only viable alternative. The savings represented will be about $10M.*

Details of the plan and supporting analysis are provided in the following sections.

Overview

LLNL is developing an Institutional 25 Year Plan for facilities. Livermore Computing (LC) has been tasked to prepare a Long Range Facility Master Plan for the TSF Complex facilities and infrastructure that includes buildings B-115, B-117, B-439, B-451, and B-453. Although the Master Plan considers requirements from all programs potentially requiring HPC and data intensive computing facilities, this brief focuses on the facilities and infrastructure of relevance to LC and the ASC Program. The Master Plan assesses existing facilities while considering interim (5-10 year) and long-range (up to 25 year) goals for HPC and data intensive computing at LLNL. The plan identifies deficiencies and addresses opportunities for improvement and modernization. Based on an analysis of the existing facilities and infrastructure, considerations have been made relative to reuse, refurbishment, and replacement of existing facilities. The Master Plan:

- Includes sustainability and energy efficiencies as an integral part of future facilities
- Evaluates existing facilities and site-wide infrastructure
- Incorporates current processes and future plans that might influence the Master Plan effort
- Evaluates future trends for required computing needs
- Makes recommendations for Deactivation and Decommissioning (D&D), renovating or reusing existing facilities, and identifies sites and program components for potential future facilities.

Strategy

Gap Analysis

An HPC gap analysis was performed to evaluate existing facilities and infrastructure for current and potential future capabilities for HPC applications. Figure 1 depicts the current LC-managed HPC facilities.

Building	Area (ft^2)	Classification	PUE	Year Constructed	Installed HPC Power Capacity (kW)
B-115	4,576	U/C	1.82	1954	500
B-117	5,418	C	1.43	1953	750
B-439	6,364	U	1.82	1980	1,200
B-451	19,000	U	1.42	1980	3,600
B-453 Main	48,000	U/C	1.30	2004	30,000
B-453 Aux Rms	4,885	U/C	1.49	2004	725

Figure 1. LLNL HPC facilities and capabilities

A process was developed to assess the probability of siting future extreme scale systems and infrastructure to meet mission requirements in existing facilities. LLNL's facilities were evaluated for structural (S), electrical (E), and mechanical (M) capabilities. A ranking system was created to assist with this evaluation based on the following system component breakdown:

- Advanced Development Systems
- Commodity Systems
- File Systems
- Storage
- Network/Support Systems

A *High* rank indicates that the infrastructure will probably be acceptable for future computational systems without modifications or upgrades. A *Medium* rank indicates that the infrastructure is probably acceptable for future computational systems with minor modifications and upgrades. A *Low* rank indicates that the infrastructure is likely to be unacceptable for future computational systems without significant modifications and upgrades. A *Fail* rank indicates that the infrastructure is unacceptable for future computational systems and, therefore, the mission cannot be achieved.

All current HPC facilities were evaluated against these criteria and the results are depicted in Figure 2. These results are for 15-year time horizon. It should be noted that there are immediate issues such as that B-115 and B-117 are unable to meet mission today with respect to advanced development and commodity systems. That is not the case at this time for B-453 but, projecting out 15 years, significant infrastructure upgrades are likely to be required.

	B115		B117		B439		B451		B453	
Advanced Development Systems	SEM		SEM		SEM		SEM		SEM	
Commodity Systems	SEM		SEM		SEM		SEM		SEM	
File Systems	SEM		SEM		SEM		SEM		S	EM
Storage	EM	S	EM	S	EM	S	EM	S	SEM	
Network/Support Systems	EM	S	EM	S	EM	S	EM	S	SEM	

Color Legend:		
	High	No Upgrades Required
	Med	Upgrades Required
	Low	Upgrades Cost Intensive
	Fail	Unable to Meet Mission

Coding Legend:	
S	Structural
E	Electrical
M	Mechanical

Figure 2. Evaluation and ranking of HPC facilities – 15 year time horizon

As shown, all existing facilities except for B-453 fail to meet needed future (short-term or interim) structural, electrical, or mechanical requirements for three of the five major components (in particular, the supercomputers themselves, whether advanced development or commodity systems). Even B-453 is marginal in that significant upgrades will be needed to site future advanced development or commodity systems. Existing facilities meet requirements only for relatively low-density, low-power equipment. The oldest facilities (B-115 and B-117) require major seismic upgrades; this is not unexpected, considering their age—nearly 60 years old.

Master Plan Focus

Because future HPC systems will push beyond many of the capabilities of the existing facilities, including their electrical, mechanical, and structural infrastructure, the focus of the Master Plan is to:

- Utilize existing facilities based on their capabilities and perform reasonable upgrades and renovations
- Demolish aging facilities that can no longer meet mission requirements
- Build new facilities to meet mission requirements
    - Focus on sustainability and energy efficiency through flexibility
    - Scale footprint with the computational technology

- Deploy innovative HPC facility design methodologies

Core Competencies

LLNL has spent the past several years focusing on sustainability in HPC and has identified core capabilities or competencies needed to improve operational efficiencies and ultimately achieve extreme scale computing:

- Benchmarking
- Computation Fluid Dynamics (CFD)
- LEED Certifications
- Leverage Existing HPC Capabilities
- Free Cooling
- Liquid Cooling Advances
- Innovative Electrical Distribution
- Facility Power Management
- HPC Facility Gap Analysis
- Sustainable HPC Modular Facility

The core capability of the *Sustainable HPC Modular Facility* is key to this Master Plan. The new facility concept is *to construct facilities using a scalable building concept that meets changing HPC demands by scaling the power, square footage, structure, and cooling solutions in modules and only when needed*. These new modules will deploy facility power management systems, minimize the use of cooling towers and chillers, and utilize closed-loop cooling concepts with expanded liquid cooling temperature ranges. This concept is more cost effective and allows LLNL to leverage existing HPC facility practices. Illustration of a single, double and triple modules are shown below along with cross section.

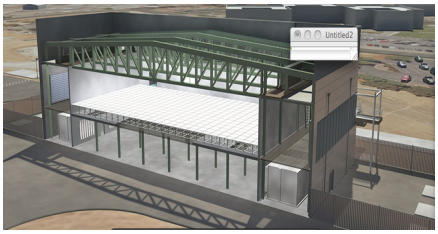

The primary rational for a new scalable facility is driven by the (1) extreme cost of renovating B-451 so that it can support near term future systems, (2) and the inability to continue to upgrade B451 into the indefinite future in a sustainable and scalable manner. Upgrades to B-451 necessary to support unclassified systems over the coming decade and associated project costs are shown in Figure 3. Indeed, $20.5M must be invested over the next two years to upgrade the facility whereas only $10M would be required to build the first 6,000SF and 6MW of a new Sus*tainable HPC Modular Facility* that can grow and adapt as needed as requirements are identified in the future.

Location	Project Title	Project Scope	FY13 Cost ($M)	FY13 Duration	FY14 Cost ($M)	FY14 Duration	Total Costs ($M)
B-451	B-451 3MW Open Computing Power Upgrade	Increase power for Open Computing in the existing low density areas of the facility. Provide an additional 3MW of computational power.	$4.0	12 months	$4.0	12 months	$8.0
B-451	B-451 Central Chilled Water Plant Improvements	Provide additional capacity for future Open Computing commodity and advanced development systems.	$2.0	12 months			$2.0
B-451	B-451 Liquid Cooling Tertiary Loop	Provide liquid cooling solution for future Open Computing commodity and advanced development systems to the raised floor.	$2.0	12 months	$1.0	4 months	$3.0
B-451	B-451 Structural Modifications	Provide adequate structural improvements to the computer room to accommodate future systems and bring facility up to current codes. Original facility was constructed in 1980.	$4.0	12 months	$2.0	8 months	$6.0
B-451	B-451 Network Infrastructure	Provide adequate network infrastructure to accommodate Open Computing in the facility			$1.5	8 months	$1.5
Total			**$12.0**		**$8.5**		**$20.5**

Figure 3. Building 451 upgrades to support future systems

Fundamental Conclusions

As has been discussed above, leveraging existing facilities, renovating existing facilities, demolishing aging infrastructure and building new facilities are all options that have all been considered to meet future unclassified and classified computing needs.

There are at this time two high level conclusions regarding facilities housing existing and future ASC program infrastructure:

- B-453 will be fully able to serve as the primary residence of classified HPC operations for the next decade at least. However, a plan must be developed to assure that B-453 is ready to house the FY17 Advanced Technology System as well as capacity clusters sited in that time frame. This will be the subject of an analysis and white paper to appear in CY13, once the AT system contract is in place and the Program has a detailed

understanding of the requirements coming from that AT system in particular. This will assure the most effective and least costly solution.

- LC will be unable to site expected unclassified FY14 TLCC3 systems, just as it was unable to site Vulcan, without an investment in infrastructure. By far the optimal alternative, when considering both cost and a long-term strategy is to construct the first module of a *Sustainable HPC Modular Facility,* with work beginning in FY13 and completing in FY14, in time for siting unclassified TLCC3 scalable units.

Open computing is critical to the ASC program. However, it also is potentially a future business opportunity for LLNL. ASC's mid-term needs will be fully met with the addition of a single module as long as there will be supplementary support from B451 for supporting infrastructure (and B453 with Vulcan). Anticipating future needs coming from ASC and new business, architectural plans have been developed which would allow for the construction of additional modules, but *only when these needs are pressing and clearly identified*. This growth could potentially require up to 15MW and 24,000SF of computational space over the next 25 years. As these opportunities evolve, the *Sustainable HPC Modular Facility* strategy will be employed to meet these growing demands – each module will employ the best technology available at that time for cooling and support. The *Sustainable HPC Modular Facility* will be located in proximity of utilities and positioned so it can be renovated or expanded at low cost to meet future unclassified HPC demands.

Schedule and Migration Plan

Figure 4 provides the construction schedule for the proposed *Sustainable HPC Modular Facility*. Design is complete. Upon funding approval, a bid package will go out, the contract will be awarded, construction will ensue. Commissioning is projected for FY14Q3 after which the facility will be available for siting new equipment, TLCC3 in particular.

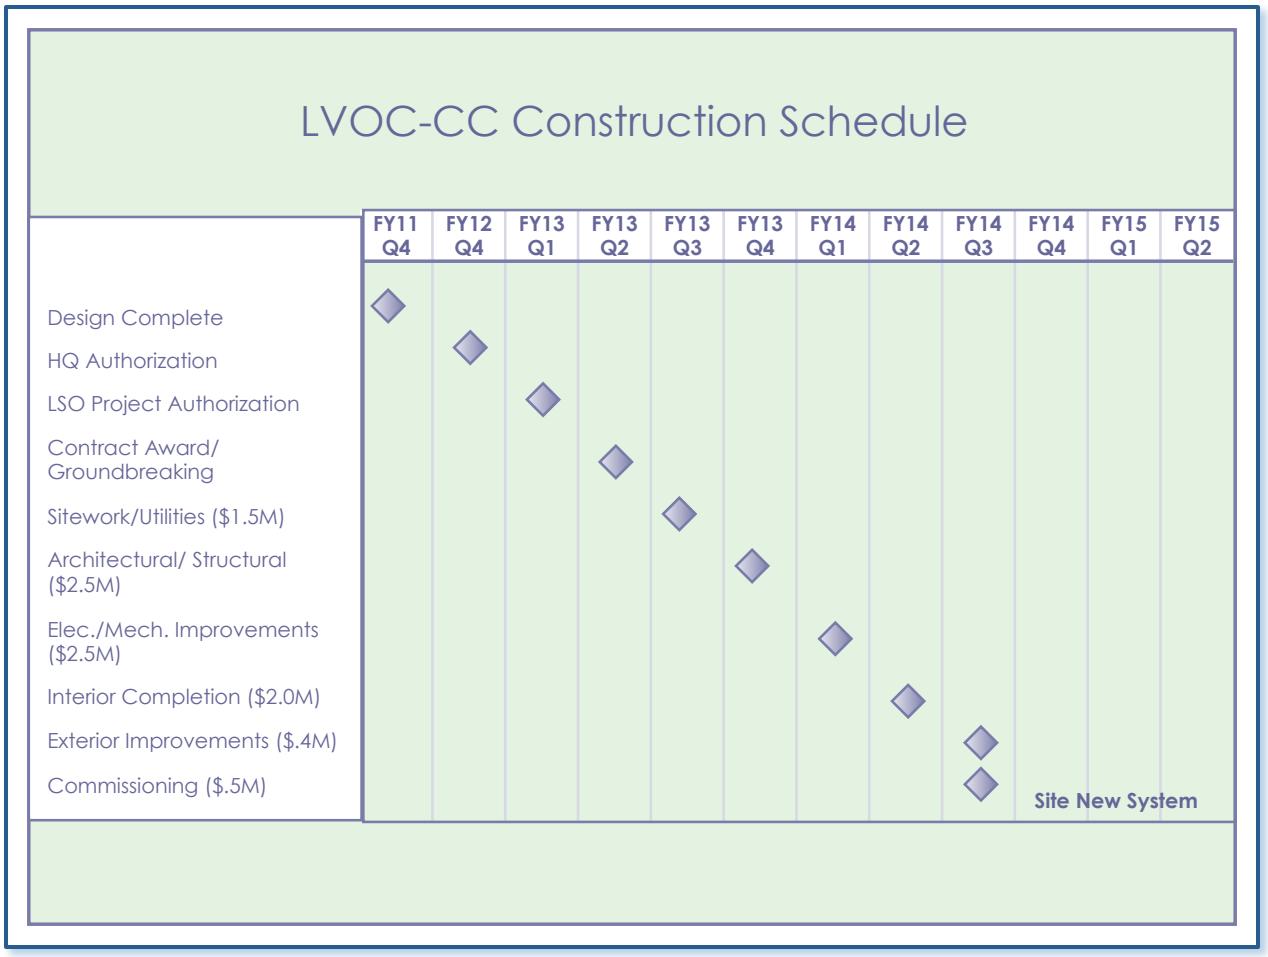

Figure 4. Sustainable HPC Modular Facility construction schedule

Completion of the *Sustainable HPC Modular Facility* will set in place a migration plan. The high level details of this plan are to move all systems out of B-115 and B-117 allowing for their D&D. The current use of existing facilities by component type is depicted in Figure 5. Displaced equipment from B-115/B-117 will be targeted for either B-439 or B-451, one of which[2] will need to be classified in order to support the classified infrastructure coming out of B-115. The new *Sustainable HPC Modular Facility* will be targeted for new systems such as TLCC3. In the long-term, B-439 and B-451 will continue to support storage and network/support systems. The migration plan will most likely take two to three years to complete and would also allow for some systems to be retired in place rather than go through the expense and risk of moving them.

[2] Most likely B-439.

Building	Advanced Development Systems	Commodity Systems	File Systems	Storage	Network/Support Systems
115				C/U	C/U
117		C*	C*		C
439		U*			U
451		U	U	U	U
453 Main	C/U	C	C/U	C	
453 Aux Rms					C/U

Coding:
C Classified
U Unclassified

* Very limited use currently and not extensible

Figure 5. Component usage of existing facilities

Request of ASC HQ

*LLNL requests that the HQ NNSA ASC program concur with execution of the first module of this effort. The LLNL ASC program suggests that $5M FY13 and $5M FY14 of LC FOUS infostructure operating funds be used for a GPP project to complete the module in FY14. LLNL **does not** request additional funding from HQ for this work. The 5-year LLNL infostructure plan has allocated the necessary funding, as this strategy is far less expensive than retrofitting B451, the only viable alternative. The savings represented will be about $10M.*